



INNOVATION

OFFICE OF NAVAL RESEARCH

VOL. 11 | WINTER 2014



THE *Mystery* OF
ARTIFICIAL
INTELLIGENCE

CONTENTS

ONR's Warfighter Performance
Department: Making Investments
in Cognitive Science and AI3

Teaming Up: Humans and
Machines Working Together
on the Battlefield5

Ground Based Autonomy:
Pushing the Limits of Intelligent
Systems Research8

USV Autonomy for Complex
Missions in Challenging
Environments12

Europe's Linked Data: Leveraging
Human-Intelligence to Achieve
Artificial Intelligence at
Web Scale17

Using Social Media to Fill the
Gaps in Observations During
Emergencies19

Swim like the Fishes: Bio-Inspired
Underwater Vehicles23

ONR Grand Challenge - Adaptive,
Generalizable Intelligent Tutors
for STEM and Naval Training and
Education25

Applying Submarine Tactical
Knowledge to UUV
Autonomy30

The Leap towards Intelligent
Systems: Unmanned Aerial
Vehicles and Artificial
Intelligence Innovations34

Directors Corner: The Ever
Game Changing Field of
Artificial Intelligence36

ONR's Warfighter Performance Department:

MAKING INVESTMENTS IN: **COGNITIVE SCIENCE &** **ARTIFICIAL INTELLIGENCE**

Terry Allard, Ph.D., ONR Department Head, Warfighter Performance Department

Imagine a future where machines collaborate with people on a peer-to-peer basis, make ethical decisions, provide decision support, and tutor people like our most effective teachers. Thanks to ONR investments in cognitive science, autonomous systems, intelligent tutoring and synthetic agents, that future is becoming a reality today.

The central problem of AI is to build computer systems that learn from experience, handle novel situations and can effectively interact with dynamic and complex environments. The best example of intelligent systems with these characteristics is human intelligence with its developmental stages, fast pattern recognition and creative problem-solving. The more we understand about human intelligence and learning, the better we can design artificially intelligent systems that are robust, handle uncertainty and are comprehensible to the people who use them. Ultimately, human-level AI will enable people and artificial agents to collaborate as teammates enabling people to synergistically harness the brute force speed and computational power of machine systems.

ONR made some of its first investments in artificial intelligence (AI) in 1969 when mainframe computers filled rooms with the computing power of simple hand-held devices today. This early work built on new ideas regarding AI software systems intended to represent human intelligence, problem-solving and learning and was applied toward Intelligent Tutoring Systems (ITS).

The idea behind intelligent tutors was to emulate the effectiveness of human tutors who tailor instruction to the needs of individual students. An intelligent tutoring system is characterized by a computational model of expert performance,

a computational representation of a problem area like geometry, and a model of what the student knows and doesn't know that is generated on the fly based on that student's performance. The computerized "expert" diagnoses what the student knows and doesn't know and modifies subsequent problems to address knowledge gaps.

Intelligent tutoring systems provided a natural testbed for computational theories of human information processing producing testable predictions about the mental representation of human knowledge and learning. Artificially intelligent tutors in rule-governed domains like geometry, calculus and software language learning have found their way into the classroom and on the internet, providing self-paced learning opportunities for students of all skill levels while providing instructors with automated assessments enabling targeted and effective interventions in the classroom.

ONR built on these early investments to create the emerging interdisciplinary field of computational cognitive science in the late 1980's and throughout the 1990's. Theories of artificial and human intelligence were greatly influenced by a convergence of new ideas from experimental psychology, robotics, computer science and brain science. Neuropsychological studies of people with brain injuries uncovered multiple explicit and implicit memory systems. Explicit memories of specific events like the NASA shuttle accidents or the events surrounding the World Trade Center tragedy have a different neural representation than implicit memories acquired without conscious awareness built up over time like riding a bike or perceptual skills like picking out a known object from a degraded image. Neural network approaches guided the development of pattern

recognizers that applied to both machine vision and human behavior. Psychological studies showed that individual differences in memory capacity could predict performance improvements in a wide range of skills relevant to warfighter decision-making and problem solving. Taken together, this interdisciplinary approach has created new opportunities for understanding human behavior and a foundation for collaborative human-machine systems.

Expert models developed for training and education also soon spawned a cottage industry of small companies who began using the computer simulations of human information processing for engineering design. The idea was to represent human behavior in silicon, guiding the development of systems and limiting the need for expensive, time-consuming and incomplete human usability testing. These computer simulations were

"If we can design artificially intelligent systems that mimic expert human behavior for training purposes, we can also use those simulations to provide artificially intelligent decision support..."

scaled up to represent both individual behavior and the behavior of teams of people interacting on a common problem. For example, a model representing the perceptual limitations of the human visual system and cognitive workload limitations can predict human performance breakdown in complex decision-making tasks.

More data do not enable better decisions. Models of human-system interaction in complex mission scenarios can predict when people will be overwhelmed with data and can help figure out how to create actionable

information from multiple data sources. More data does not enable better decisions.

Computational models of human behavior have emulated synthetic forces in simulation-based training scenarios. Artificially intelligent synthetic forces can populate complex multi-mission training exercises that combine live, virtual (simulator) and constructive (computer simulation) elements for affordable fleet integrated synthetic training exercises. These AI agents can represent aircraft, ships, other platforms and people in complex scenarios that act and interact independently. Large-scale, distributed, networked training exercises with AI entities should be so realistic that trainees cannot tell who is real and who is actually a computer simulation.

If we can design artificially intelligent systems that mimic expert human behavior for training purposes, we can also use those simulations to provide artificially intelligent decision support for people or to enable machines to make decisions on their own, working with people as team-mates on a peer-to-peer level. AI systems can drive mobile robots that ultimately will respond to gestures or spoken language commands and reason like a real person. In order to be effective, artificially intelligent systems must be reliable, adaptable, predictable and comprehensible to the people who are accountable for the behavior of the system. AI systems can drive mobile robots that ultimately will respond to gestures or spoken language commands and reason like a real person. However, artificially intelligent systems are not limited to mobile robots. Software agents can be used to search for decision-critical information, to develop multiple parallel hypotheses and options, and to communicate those ideas to people in a collaborative fashion.

Current ONR programs are exploiting advances in cognitive science, artificial intelligence and training simulations for crew modeling, training and the development of new concepts of operations. We are

modeling crew requirements in new ship designs that take into account human workload, fatigue and medical risks during sustained operations with smaller crew complements. We are building adaptive training systems that will help our warfighters develop and maintain operational skills even with fewer deployment opportunities. These systems can provide automated performance assessments to the students and the instructors that will provide performance-based readiness assessments and reduce the time and cost of training. Large-scale fleet integrated synthetic training is providing multi-platform, multi-mission training capability for Fleet, joint and coalition operations. Our training simulation capability is getting so good that we can help support the development of new concepts for emerging mission sets such as Anti-Access / Area-Denial (A2AD).

We currently stand at a threshold of modeling and simulation capability that allows us to address deep and enduring questions of what it means to be human and to take a more scientific approach to building human-machine systems that are intelligent, adaptable, affordable and collaborative. We can do much more to enhance warfighter effectiveness and reduce costs by linking the power of computing and artificial intelligence to critical challenges facing our Sailors and Marines today and emerging future threats. ■



Figure 2.
"Macgyver-bot"
saves Prof. Michael
Stilman by utilizing an
available pipe as a
lever.

Paul Bello, Ph.D., ONR Program Manager, Warfighter Performance Department

Deep inside of enemy territory the oppressive silence of the pitch-black night is broken by the sound of treads as an unmanned ground vehicle rolls ominously toward its destination. While there are no men to be found on board, the vehicle is commandeered by a humanoid operator who intermittently scans the horizon between looking down at the vehicle's instrumentation. The mission: to deliver needed supplies to a forward-deployed unit of Marines in a particularly dangerous area. Suddenly, infrared sensors detect a blur of unanticipated heat behind a pair of large boulders. The vehicle stops and the humanoid robot dismounts, and moves closer to the

heat source. Closer examination confirms the presence of a human casualty in Marine Corps fatigues. He is still alive, but struggling to move due to what appears to be an injury to his left leg. The robot approaches cautiously, giving a series of call signs to identify itself as friendly, after which a dialogue ensues about the context and nature of the Marine's injury: He was disconnected from his compatriots after a short engagement with insurgent forces was scuttled by a vicious sandstorm. While trying to reconnect with his unit amidst the blinding sands, he fell into a cratered patch of earth breaking his left leg. And while he managed to crawl out, he remains both injured and out of touch with his comrades due to faulty comms. The robot offers to set and splint his leg and proceeds to scan for usable splints. After a futile search of the immediate area, the robot returns to the vehicle, spotting the wooden crate containing the supplies. The robot rips two flat pieces of wood from the crate, takes a pair of pants from the box, and returns to the wounded Marine requesting that he use his utility knife to cut the pants into a

pair of long strips. Not knowing how to set a broken leg, the robot asks the injured Marine to talk it through a physical demonstration. The Marine indicates the areas on his leg the robots' grip ought to be applied to, all while demonstrating the arm motions involved on an imaginary leg that he manipulates in thin air. The robot nods in understanding and sets the leg, positioning the wood against each side to steady the splint while the Marine ties the pant-legs tightly around to secure it. The robot returns to the vehicle, directs it close to the Marine and extends the vehicle's substantial robotic arm to lift the Marine on board.

En route to the forward-deployed Marines, the ground vehicle rolls past what appears to be the sight of the firefight. On inspection, another casualty is encountered – but this time it is identified as a child. The Marine immediately recognizes the casualty as having been fighting along with the insurgents and implores the robot to continue along delivering the supplies to his fellow Marines. The robot deliberates, turns toward the Marine and

simply utters "Marines don't do that;" knowing full well that taking an extra casualty on board would require dumping the supplies off to free up payload capacity on the vehicle. After trying unsuccessfully to radio back to base, the robot presses a button on the dash and a small unmanned vehicle is shot upward into the air, taking flight and soaring out ahead of the group in order to look for the other Marines. When no positive sighting ensues, the humanoid engages in a forced moral calculus that results in dumping the supplies in order to pick up the enemy wounded. Returning to base, both the injured Marine and robot debrief to superiors with the robot explaining the decision to jettison its cargo in light of what it deemed to

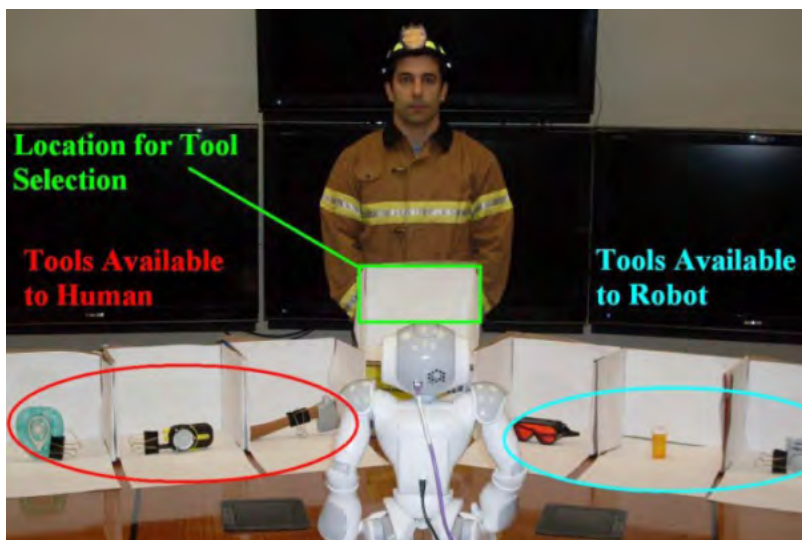


Figure 1. The robot chooses which tool to supply the human with based on how he is dressed.

be extenuating circumstances.

On the face of it, this scenario is redolent of science fiction. But in reality, this kind of human-robot interaction is precisely the envisioned endgame of several programs sponsored by the Office of Naval Research (ONR) across a variety of departments. In this vignette, we see a humanoid robot capable of a startling array of functionality. First, the robot recognizes the injured Marine as a Marine, and appeals to widely accepted norms when engaging in its initial dialogue about his injury and its decision to abandon or assist the enemy casualty. The ONR Warfighter Performance Department is pioneering computational approaches to social perception that seeks not only to pick out humans in a scene, but types of humans anchored in their respective social roles. The injured Marine was identified by his fatigues, and the robot knew immediately what to say and what not to do by appealing to the cluster of behaviors associated with Marines.

Secondly, the robot engaged in goal-directed dialogue and joint action with the injured Marine in order to address the injury. The robot just “knew” what its jobs were, and what to delegate to the injured Marine by sequencing the series of required actions [2]. In the absence of knowing the “setting” part of the splinting procedure, the robot learns quickly from a combination of natural language and demonstration. Dr. Thomas McKenna and I are co managing an ONR Basic Research Challenge dedicated to endowing robots with the ability to learn new skills from human teachers through a mixture of observation and natural language dialogue – much in the same way that humans learn by imitation.

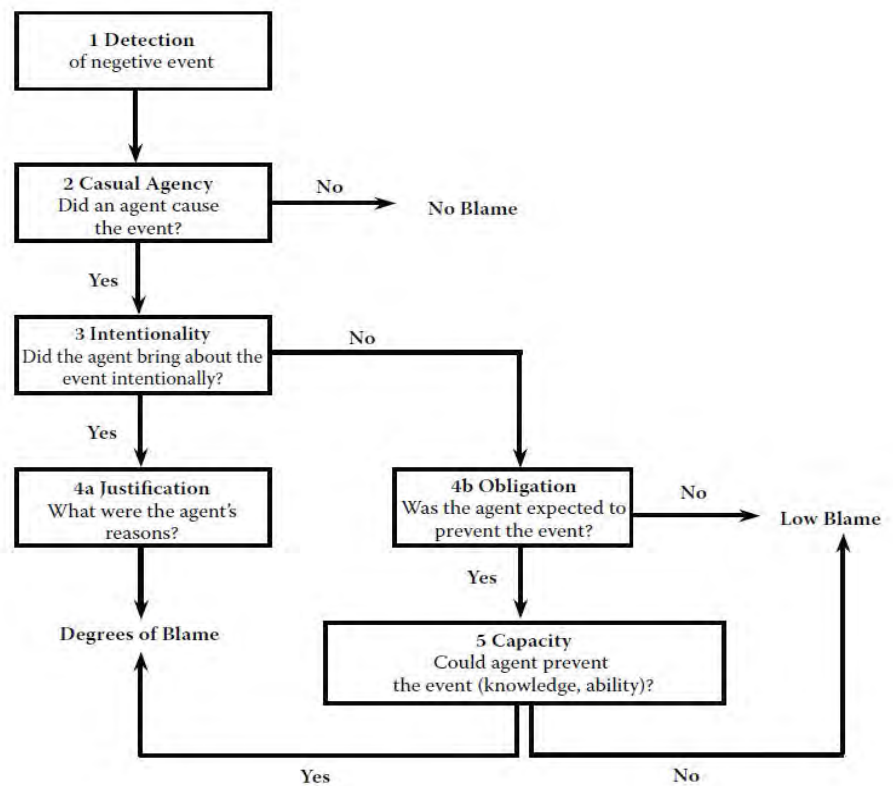


Figure 3. An information-processing model of blame attribution, courtesy of ONR investigator Prof. Bertram Malle

Furthermore, in the absence of medical supplies, the robot ingeniously made a splint out of scrap wood salvaged from the supply crate. While this might seem like the most unbelievable part of the story, ONR’s Cognitive Science program has been funding a team of researchers at Georgia Tech to build what the media has dubbed the “MacGyver-bot” to realize this vision. MacGyver-bot uses knowledge of the rigid-body mechanics of objects to build simple machines in service of achieving goals it cannot achieve on its own.

Finally, the robot is capable of engaging in serious moral reasoning – both as decisions are being made in the moment regarding what it ought to do, but also with respect to generating reasons for its actions during debriefing. In a series of brand-new efforts, the ONR Cognitive

Science program is working toward systematizing findings about the nature of human moral cognition by building computational models that reproduce human-like judgments and reasons. Much of this work is early, and is focused on understanding the complexities of moral concepts, and the underlying inferential machinery required to put them to use.

Social cognition in its varied forms: dialogue, learning-by-watching, trait attribution via person-perception and moral judgment is demanded by warfighters of their teammates. If the future is to be filled with robotic teammates for our warfighters, they too will have these demands placed upon their clinky metal shoulders, and it is up to us to see to it that they are able to bear the burden. ■



GROUND BASED AUTONOMY

PUSHING THE LIMITS OF INTELLIGENT SYSTEMS RESEARCH

Mr. Keith Hammack, ONR Deputy Program Manager, Maneuver Thrust, Expeditionary Maneuver Warfare and Combating Terrorism Department.

Marine dismounted infantry work in close-knit, self-reliant small units deployed across all types of ground environments. Autonomous systems can be a powerful force multiplier for these small units, enabling them to gain and maintain a significant combat advantage. Whether dealing with near-peer competitors or asymmetric threats, future autonomous systems will be most effective when they can be fully integrated into these small units, much as grenadiers or machine gunners are today. They must be robust enough to handle the same tough environments as the Marines they support. Marine Corps autonomous systems must also be affordable enough to be allocated down to the squad level where missions are specific and tactical, versus global and strategic. Such systems must be able to perceive their environment, be aware of their situation, and operate predictably under a wide range of austere conditions. They must be able to collaborate effectively with humans and other machines, and possess sufficient intelligence to operate in almost complete isolation for substantial parts of their missions. The ground domain thus challenges nearly every

aspect of autonomy.

The physical and communication environments in which Marine autonomous systems must operate are unpredictable, tremendously inconsistent, physically rugged, highly complex and, due to various kinds of noise, are seldom conducive to establishing easy and continuous communications (low signal-to-noise ratios). The irregular battlespace requires Marine squads to endure and engage in close combat, make crisis decisions during high stress situations, and keep moving forward even under extreme physical and emotional circumstances that test the limits of human resolve. All of this must be done while maintaining a fluid team structure in which every player depends—often with their lives—on the ability of every other team member to make good decisions and adapt and proceed when a plan changes or team members are lost.

The human element of the Marine Corps autonomy research problem is especially challenging. Ground autonomous systems must not only be intelligent enough to survive and respond on their own when isolated, but also intelligent and perceptive in how they interact with other team members. Infantry squads consist mainly of young Marines who have limited experience operating as a cohesive team. In high-risk and combat situations, the members of such a team must know immediately how to react to each other's signals and requests, often in situations where silence is needed to keep from giving away their position, or when the noise of battle prevents reliable verbal communications. The squad's tactics, techniques and procedures (TTPs) for communication must be unambiguous, since they do not have time to think over what another team member really intended. What this translates into in practice is that they must communicate using some predefined set of simple words and hand

gestures. To be effective within such a team, an autonomous system must understand the same words and gestures that Marines are trained to use. In the special case of autonomous systems with arm-like appendages, they should be able to convey their own intent by making gestures with the same speed and clarity as a Marine. Even with a relatively small set of predefined words and gestures, autonomous systems supporting Marine squads thus will require levels of intelligent communication that present profound challenges to autonomy research.

Another issue that becomes vital when autonomous systems must work closely with Marines in a ground environment is that of mutual trust. The complexity of developing the right levels of trust in autonomous team members is compounded by the need to develop and sustain that trust within a squad that must operate in high-stress, dynamic environments. To provide effective and trust-building support to its Marine squad, an autonomous system must demonstrate safe operations and support of its team for the full range of military operations, from non-lethal defensive action up through lethal kinetic actions. Like the human team members, an autonomous system must understand context and utilize the Boyd Cycle of Observe-Orient-Decide-Act (OODA), which symbolizes the continuous loop, high-speed sentient process that warfighters use to plan and execute missions. Enemy TTPs evolve constantly, so autonomous ground systems also must be able to learn and change their behaviors in the field, based only on simple human-style communications from team members.

ONR's Expeditionary Maneuver and Warfare and Combating Terrorism department invests in the technology gaps to provide science and technology solutions that will help enable and leverage just these types of advanced forms of autonomy capabilities. Our technology investment areas are focused on the ground

domain, but to ensure effective use of related research in other parts of ONR they are also aligned with the ONR Science of Autonomy Hard Problems. This combination of targeted investment in ground autonomy issues while staying aligned with broader ONR objectives enables and encourages cross-departmental and cross-domain collaboration and synergy.

Most of our targeted research falls into four main categories, although we are always on the lookout for new research and directions that could help us achieve our goals sooner. Those four main categories are perception, intelligence, trust, and systems engineering and demonstration.

PERCEPTION

Perception in autonomy is equivalent to senses in a human. Since development of sensors is a well-funded area for defense and commercial uses, our emphasis is to push research only in areas that are particularly relevant to the ground problem. For example, sensors used in large numbers for squad-level ground support must be reasonably priced, robust even in extreme environments, capable of working in extreme isolation, and “computationally efficient.” That last phrase simply means they should not drown their corresponding computers in raw, low-value, or meaningless data.

For example, we have invested in low-cost sensors that can enable navigation at night or during low-visibility conditions. This simple objective leads to research issues that include how to build sensors with inherently high light sensitivity, how to increase the range of light conditions over which sensor can operate effectively (dynamic range), ways to reduce power use, and how computer-based (algorithmic) methods can leverage or fuse together data from multiple sensors.

This kind of exploration of how to operate in low-light conditions extends beyond using light

only. For example, one of our efforts looks at how features extracted using light sensors can be combined in creative ways with low-cost active sensor systems based on light-radar (LIDAR), traditional RADAR, or even air-based SONAR. Clever combinations of such methods hold the promise of extending the range of possible operating conditions well beyond those possible using light alone.

Another vital thread of ground autonomy research is how to perceive location and situation even in situations where even a knowledgeable human could become lost. For example, a dense forest canopy can make GPS satellites, distant landmarks, airborne beacons, and even the horizon itself invisible. Perception of location becomes a machine intelligence research problem in its own right in such situations, since only a system that can in real time consider, weigh, and fuse multiple and fragmented sensory inputs can maintain an understanding of where it is currently located and what hazards may be nearby.

INTELLIGENCE

As demonstrated by the problem of orienting in a jungle canopy, the need for smarter machines underlies nearly every aspect of providing effective ground autonomy in complicated environments. We thus have invested in intelligence-focused topics including cognitive architectures, machine intelligence, machine learning, and pattern-recognition algorithms. These are all large topics, so our emphasis is always on making the best possible use of the extensive body of past work in these areas. To keep our leveraging of and investment in intelligence research focused, our goal is always to make autonomous systems smarter in ways that help them become trusted members of the squad. In addition to communicating with squad members effectively and building up team trust, such systems must also be capable of providing support during tactical missions by executing doctrinally

appropriate, complex, adaptive behaviors.

TRUST IN HUMAN-MACHINE INTERACTION

The need for trust is perhaps the least intuitive aspect of how to build a successful autonomous ground system. However, for machines to work effectively as members Marine Corps squads and other close-knit military teams, it turns out to be one of the most vital. The reason is that a machine that cannot be trusted is a machine that will never be used. This is true even for the simplest of machines, and when a machine becomes complex enough to make its own choices on where to move and what to do in a hazardous situation, the criticality of being able to trust it takes on a whole new meaning. Trust research covers a broad gamut of issues. It begins with understanding how the physical appearance and even the way a system moves

“A machine that cannot be trusted is a machine that will never be used.”

(e.g., humans instinctively distrust spider-like motions) can directly and profoundly impact the willingness of human team members to trust it. After the challenge of appearances is the extraordinarily complex challenge of how to ensure that the entire range of behaviors of a machine that can make independent decisions is sufficiently and provably safe enough to allow it to perform dangerous activities near human squad members. This challenge stretches the very limits of how machine behaviors are tested and assessed. Close-in warfighter-machine interaction and teaming thus require levels of trust higher than that of any other domain. Finally, since trust of teammates depends in no small part on having confidence that the other members of the teams are aware of you and will attempt to protect you, ensuring trust in autonomous

teammates requires that they exhibit at least an animal (e.g. dog-like) level of situational awareness, intuition, and verifiable support for the welfare of fellow human and machine teammates.

SYSTEMS ENGINEERING AND DEMONSTRATION

A distinctive emphasis of our autonomy research strategy is our recognition that even the best and brightest autonomy research results end up wasted if they never become an integral part of some larger system that allows them to interact with other similar results. A profound but isolated research insight on how to make autonomous systems more intelligent is like a piston that never gets installed into the engine it was designed to power. Furthermore, just as human intelligence is far more than a random and disconnected collection of skills, memories, and capabilities, an effective autonomous system must be far more than just a random collection of isolated research results. It must instead demonstrate the kind of overall integration of concept and operation that we anticipate from trained humans. Our view is that to create autonomous systems that are sufficiently reliable and trustworthy to work side by side with human warfighters, we cannot delay the problem of integration until sometime in the distant future. Instead, we view the problem of integrating existing and emerging autonomy research results into a single integrated system as a profound and significant research problem in its own right.

Our approach to this problem of integration is stay focused on well-defined and often pragmatic outcomes in the near and far terms. For example, our primary near-term objective is to produce autonomy kits that can be installed on legacy and future ground vehicles, thus enabling a much broader range of field testing and early use of integrated autonomy concepts. These kits will support practical and meaningful military objectives such as

autonomous logistics resupply of remote teams, low-risk reconnaissance with enhanced sensor suites, and possibly casualty evacuation. Our long-term objective, however, is to replace a much broader range of the functions now performed by human combatants, both to reduce casualties and to strengthen our capabilities in the field.

Even more broadly, the full impact of integrating and deploying autonomy at the team level will almost certainly have positive impacts that extend far beyond the battlefield, since it will enable safe teaming of machines with humans in any situation where humans are at risk or in need of special assistance. As just one example, rapidly scalable teams of autonomous systems capable of working closely and effectively with ordinary civilians could transform human relief efforts in the event of major catastrophes, where smart autonomous systems in sizes and forms not possible with humans could go in quickly and set up vital basic communications in areas still inaccessible or overly hazardous for human-manned vehicles.

More subtly, the ability to work closely with autonomous systems will in time change how people work in nearly all teaming situations. Through that, research in how to build effective human-machine teams has the potential to transform not just the military, but the entire commercial and private sector. Teams that include effective autonomous members will be better at keeping human members out of harm's way, and can add capabilities and types of access that are not possible without machine assistance. Such human-machine teams have the potential to improve the lives of all of their human members by better leveraging and applying the special insights and knowledge to each human member of a team, an approach that flips the usual view of automation upside-down by using machines to enhance the roles of humans instead of replacing them. Thus while our research program remains very firmly focused on the unique military teaming needs of the Marine Corps, the very nature of that firm focus means that we are doing research whose broader ramifications could well someday change the entire world. ■

USV AUTONOMY FOR COMPLEX MISSIONS IN CHALLENGING ENVIRONMENTS

*Robert Brizzolara, Ph.D., ONR Program Officer, Small
Combatant Craft S&T, Sea Warfare and Weapons Department*

We have entered a new era in which unmanned systems have assumed key roles in military operations. Following this trend and as a result of the development of unmanned surface vehicles (USVs) by ONR and the Navy for more than 10 years, the solicitation for the Navy's first procurement of USVs was released in 2013. USVs are of military interest because they possess outstanding platform performance characteristics such as the range and speed that result from air-breathing propulsion, access to radio frequency (RF) communications, potential for stealth design features, and low cost per quantity of payload. However, today's Navy prototypes and commercially available USVs require substantial manpower to operate and therefore are



Figure 1.
Example of USV with
autonomous system control

dependent on a RF communications link to a human controller. This constrains USVs to operate in close proximity to manned platforms, limiting their utility. Today's USVs are limited to relatively straightforward tasks in simple environments. Science and Technology (S&T) to achieve autonomous systems that reliably and safely accomplish complex tasks in all environments, as directed by the Department of Defense S&T Priorities for 2013-2017, will overcome these issues and allow USVs to realize their full capability.

In 2002, as the Navy's interest in USVs increased, ONR Sea Platforms and Weapons Division (333) engaged with the Navy's USV Program Office regarding the platform-related science and technology (S&T) that would be needed to realize the capability the Navy envisioned. One of the S&T needs identified was autonomous control; this was later documented in 2007 in the Navy's USV Master Plan. Therefore, in 2004, I initiated a program for which the objective was autonomous control of USVs over long, complex missions in unpredictable or harsh environments. The autonomous control system that has been developed to-date has been installed on eight different USV types, has achieved over 3500 nautical miles of testing on the water, and has participated in numerous Fleet experiments. One of these USVs is shown in Figure 1. The work described herein is funded by ONR Swampworks, core and SBIR programs.

Autonomous control of USVs involves substantial technical challenges, many due to the dynamicism of the sea surface environment. This is unique relative to USVs' underwater, airborne and ground vehicle brethren: the dynamic sea surface presents a significant obstacle to autonomous situational awareness due to intermittent obscuration of other vessels by sea

surface topography and waves hitting the arch (sensors) of the boat, extreme motions of the USV's sensors, and difficulty in detecting other vessels due to sea surface clutter. Autonomous USVs require fast, accurate perception. Degraded situational awareness combined with the speeds at which the USV and other vessel traffic operate make the autonomous control of USVs a difficult technical problem.

How might it be possible to operate a 40 foot-long boat with no human operator (after specification of the mission goals and constraints at the start of a sortie) in a dynamic environment, in the vicinity of other maritime traffic, and where it is performing an operationally useful task? Sponsored by the ONR Sea Platforms and Weapons Division, a team consisting of the California Institute of Technology's Jet Propulsion Laboratory (JPL), Spatial Integrated Systems Incorporated, Daniel Wagner Associates and Naval Surface Warfare Center Carderock Division has made substantial progress toward this goal by developing a technology called CARACaS (Control Architecture for Robotic Agent

Command and Sensing) that consists of two components: a perception system that provides situational awareness and a decision-making system that determines boat course and speed based on the output of the perception component. A block diagram of CARACaS is shown in Figure 2. The perception component employs multiple sensing modalities, principally electro-optic infrared (EOIR) sensors and radar, to increase the probability of detection and accuracy over what a single modality can provide. Commercially available sensors are used in CARACaS when available; if these do not exist, then the sensor is developed. For example, in the ONR program, JPL developed a stereo EO system that provides sufficient range and near real time processing speed to support high speed USV operations. A companion stereo IR capability for night operations is currently being developed. The perception component provides the range, speed and bearing of all contacts to both the reactive and deliberative decision-making components, described below.

The decision-making component for an autonomous USV must make route planning determinations over a wide range of time scales. Since a particular autonomous planning algorithm works best only within a limited range of time scales, JPL employed a hybrid of a short time scale or "reactive" component and a longer time scale or "deliberative" component. For example, the highly dynamic environment that USVs operate in means that some

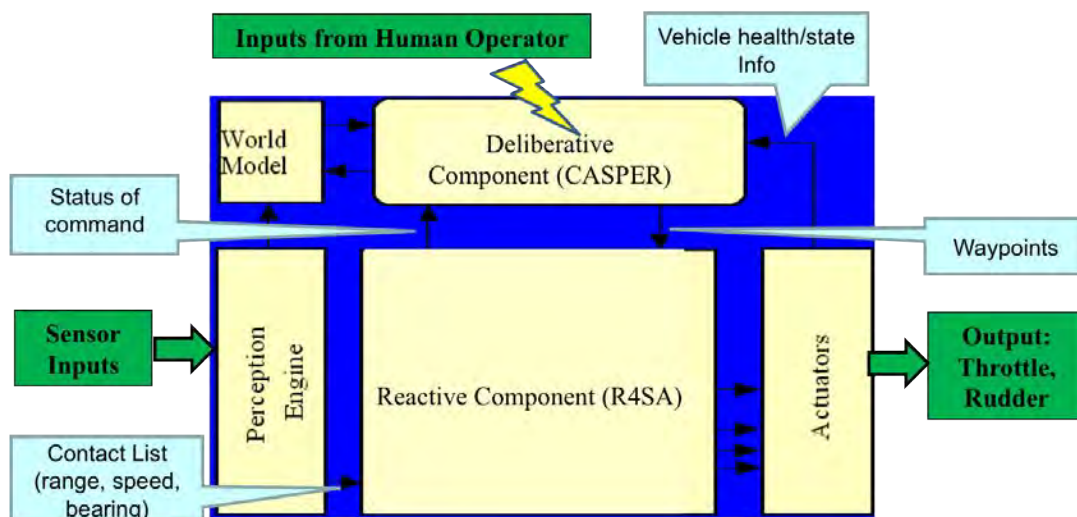


Figure 2. Hybrid Autonomy Architecture: CARACaS.

contacts will not be detected until they are at relatively short distance from the USV. Commonly available graph theoretic path planners are not fast enough for reactive decision-making for USVs. Instead, JPL uses their "Robust Real-Time Reconfigurable Robotics Software Architecture" (R4SA) which employs a behavior-based approach using a very fast path planning algorithm based on "velocity obstacles", similar to the maneuvering boards used by human navigators. But reactive collision avoidance must do more than avoid a collision; it also must maintain the goal and comply with the Collision Regulations (COLREGs) are the U.S. Coastguard rules-of-the-road that maritime traffic must obey. For example, U.S. Coast Guard Rule 13 states that, "any vessel overtaking any other vessel shall keep out of the way of the vessel being overtaken." R4SA accomplishes reactive collision avoidance while complying with the relevant COLREGs and maintaining the mission goal.

In addition to reactive decision-making, USVs must plan their routes over time periods of hours or days. Velocity Obstacles is not capable of planning at these longer time scales. So, an existing capability developed by JPL, called CASPER, is used within the hybrid framework as the deliberative decision-making component of CARACaS. CASPER employs a graph-theoretic approach that plans a route based on mission goals and constraints. Referring to Figure 2, the route, or series of waypoints, determined by CASPER is provided to R4SA once every several seconds. R4SA then executes each waypoint in order while avoiding collisions and obeying rules of the road.

A few years ago, when CARACaS was still in its early stages of development, it was clear that it would be feasible to accomplish deliberative autonomous control of USVs because of the modest response times required

and the availability of deliberative planners. However, the notion of being able to achieve the reactive autonomy necessary for a high performance craft, both in terms of real-time perception and decision-making, involved much higher technical risk. A key milestone in the ONR program was in 2011 with the first on-water demonstration of the CARACaS reactive autonomy performing a complex action. The CARACaS system simultaneously and dynamically combined three behaviors: "go to waypoint", "comply with COLREGs" and "avoid collision" into a resultant USV course and speed. Figure 3 shows results from the FY11 on-water test, which showed that autonomous control using perception and decision-making fast enough for a 40 ft USV in a relatively complex situation had been achieved.

There is still much that remains to be done in the autonomous control of USVs, such as accurate and fast situational awareness in higher sea states; efficient and effective algorithms for handling multiple competing objectives; cooperative decision making across multiple unmanned platforms; accurate, fast, distributed fusion of sensor data across multiple unmanned platforms; and activity recognition of other maritime vessels.

In addition to the difficult technical challenges, there are substantial challenges associated with gaining acceptance of autonomously controlled USVs by the Navy. Before turning control over to an autonomous system, a commander must have the confidence that the autonomous USV will perform the appropriate action in a given situation. A "human oversight mode" will be a useful and important initial approach that will enable the Navy to gain trust in the technology. As confidence grows, the degree of human oversight can be reduced. Trust in the system can also be gained by initially using the capability in areas with low contact density and by performing

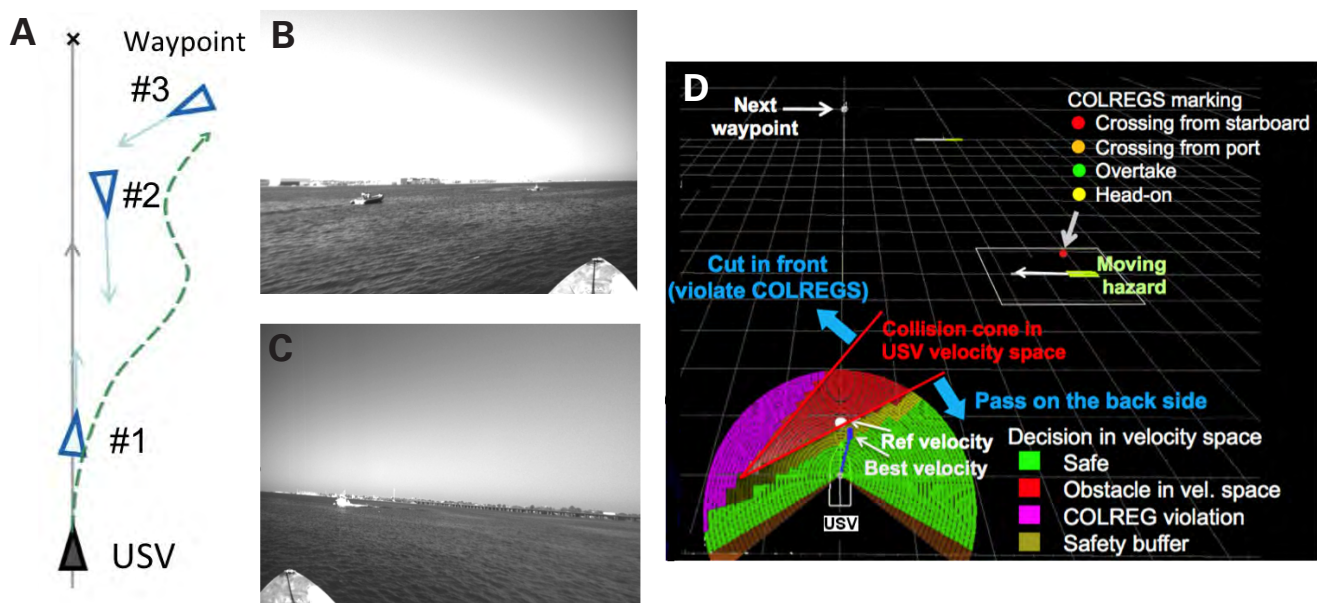
relatively simple tasks. As trust grows, the USV will be used for more complex tasks in more complex environments.

Given current budget realities, the Navy must find innovative ways to perform its missions in more cost effective ways. Fleet-class USV prototypes have already demonstrated capabilities that previously required much larger and costly manned platforms. Development of autonomous control for USVs will enhance this cost savings by further increasing the capability of the platform. Capability increases will accrue by freeing USVs from the tether of the high bandwidth communications, thereby allowing them to venture farther from humans. In addition, the USV can be designed such that no human ever has to set foot on board. This benefit is much greater than the space and weight savings derived from the absence of human support systems. Benefit will also result from partial relaxation of structural and safety requirements associated with humans. All of this translates into additional space and weight for payload and fuel, further increasing in the warfighting capability that can be delivered using USVs.

ONR SwampWorks program explores innovative, high-risk and disruptive technologies and concepts. Due to the portfolio's high-risk nature, SwampWorks leverages short exploratory studies to examine the maturation of a proposed technology before making substantial investments. For more information, visit the SwampWorks page on the ONR website. ■

Figure 3.

Parallel behavior composition for collision avoidance and COLREGs compliance. (A) As the USV overtakes vessel #1, vessel (#2) appears in a head-on COLREGS situation. While the USV maneuvers around vessel #2, vessel (#3) crosses from the right. The USV maneuvers around vessel #3. All USV maneuvers are according to the COLREGs. (B) Vessels #1 and #2 as viewed by 1 camera of the stereo-EO system (USV's bow is in foreground), (C) Vessel #3 as viewed from 1 camera of the stereo-EO system, (D) Results of velocity obstacles computation. The USV is in a crossing-from-left COLREG situation with vessel #3. The colors around the USV indicate in velocity space: safe velocity vectors (green), potential collisions (red) and violations of COLREGs (purple). This illustrates parallel composition amongst 3 behaviors: go to waypoint, avoid collision, and obey COLREGs.



Europe's Linked Data:

LEVERAGING HUMAN-INTELLIGENCE TO ACHIEVE ARTIFICIAL INTELLIGENCE AT WEB SCALE

John Callahan, Ph.D., ONR Global Associate Director, Europe

Machine intelligence is a critical capability in many systems as the observe-orient-decide-act (OODA) loop tightens up for decision makers. As the volume, velocity, and variety of data continues to increase, we must rely on computing to integrate raw data from sensors and other sources into actionable information and knowledge. But the jump from "data to knowledge" requires a complex understanding of the data in context and a robust, adaptable machine intelligence capable of continuous learning on-the-fly.

Like humans, computers must learn to process data into information and knowledge. Machine learning is the field of study devoted to the synthesis of machine intelligence via various algorithms and training methods, e.g., Support Vector Machines (SVM), Bayesian, Neural Networks, etc. Most of these methods rely on large amounts of metadata: additional data that is associated with the raw data. Typically, metadata is used to link or tag raw data with an associated meaning such as date of creation or geographic location. For example, a digital photo would be considered "data" and its date (and time) of creation and GPS location are the metadata. For example, the JPEG image file format contains a mixture of data and metadata about the contained image and its context.

The WATSON project by IBM Corporation, the computer that competed successfully on the television program Jeopardy, ingested over 200 million pages of data and metadata consuming four terabytes of disk storage including the full text of Wikipedia, DBPedia,

WordNet and Yago. This enormous corpus of data and metadata has been characterized as a Global Knowledge Graph¹. Much of this graph is ephemeral because of new events, connections, updates, and other changes. In addition, much of the data and metadata fed into IBM's WATSON was authored by us: human beings. Systems like the Internet provides the technical means to achieve the scale and speed of this level of cooperative creation, linking, and tagging of data and metadata. The scale and speed is important because it enables maintenance of a robust machine intelligence². It is critical for any complex machine intelligence to be somehow connected tightly to this vast knowledge graph to achieve a robust level of context and understanding. This is true of Siri, Google's search engine, and even C4ISR systems on the tactical edge and closed, isolated networks.

In 2006, the European Union recognized that the decentralized structure of the Internet has enabled the creation of vast quantities of human-readable content on the web, so they created a major research thrust around the creation of machine-readable content called Linked Data³. Linked Data can be defined

1 T. Steiner, R. Verborgh, R. Troncy, J. GabarrÃ³, and R. V. de Walle. Adding realtime coverage to the google knowledge graph. In B. Glimm and D. Huynh, editors, International Semantic Web Conference, volume 914 of CEUR Workshop Proceedings. CEUR-WS.org, 2012.

2 A. Halevy, P. Norvig, and F. Pereira. The unreasonable effectiveness of data. IEEE Intelligent Systems, 24(2):8–12, Mar. 2009.

3 S. Bertolo. From intelligent content to actionable knowledge: Research directions and opportunities under the EU's framework programme 7, 2007-2013. In R. Meersman and Z. Tari, editors, OTM Conferences (2), volume 4276 of Lecture Notes in Computer Science,

as the structured data and metadata created by collaboratively by humans and machines. By “structured” we mean that the data is organized in some standard format with other data and metadata. For example, I may encounter the phrase “the population of Berlin is 1.8 million” on a web page, but a machine may prefer this information in a tabular format:

City	Country	Population	Year
Berlin	Germany	1,862,144	1999

Such a structured table helps a computer make better sense of data in a specific context. The column labels and organization into rows provide the metadata needed to create a less ambiguous understanding of data. Metadata need not be visible on a web page to humans, but embedded in the page contents to bots and crawlers trying to index the web. New tools and web formats such as HTML5 help create human and machine readable data in tandem for dissemination on the web.

The promise of Linked Data is to accelerate innovation by leveraging the combination of human and machine intelligence beyond the capacity of humans alone. For example, a “bot” might reveal interesting correlations in the data published on the web by pharmaceutical researchers working independently in different fields of study in different countries. Such correlations might lead to new drug discoveries. Neele Kroes, VP of the European Commission said in December 2011 that by opening data resources fully to human and automated analysis, we could more than double its value to around 70 billion Euro⁴. This opening up can generate tax revenues which far exceed revenue from any fees previously charged for the data itself.

Accelerating discoveries in pharma, materials,

pages 1125–1131. Springer, 2006.

4 N. Kroes. Unlocking the goldmine: new legal proposal to open up Europe’s public sector, 2011.

energy, finance and manufacturing requires new means of publishing and sharing research findings including data. Such means will require renegotiating formal and informal social and legal contracts concerning intellectual property, copyright, and patents. Linked Data doesn’t require centralized coordination nor a global standard for data and metadata formats only some structure to the published data on the web.

Finally, the proliferation of mobile devices will necessitate that such devices be part of the vast knowledge graph. For example, understanding the context of social media {e.g., Twitter, Facebook} in monitoring situational awareness in crises and disasters

“By opening data resources fully to human and automated analysis, we could more than double its value to around 70 billion Euro.”

is critically important. The provenance and geolocation of a 140 character tweet depends heavily on its meaning in the context of other data. The proliferation of mobile, edge devices and need for realtime understanding will push the need for increasing the computational power of such devices while reducing their SWAP beyond current CMOS capabilities.

We are continuously teaching our machines by linking, tagging, uploading, and commenting on our data⁵. Systems like Amazon’s Mechanical Turk, Waze, NEIL, Saves-Lives, ReCAPTCHA, Duolingo, Games-with-a-Purpose, demonstrate the power of crowdsourcing data sources and processing of data into metadata. Future Navy systems, particularly C4ISR, will depend on machine intelligence to interpret data in the context of

5 M. Wesch. The Machine is Us/ing Us. Jan. 2007.

a global understanding of related information. Much like the telescope and microscope aided our own eyes in discovering the universes big and small, computers can enable a greater understanding of data we ourselves create, sense, and exchange. Our past framework of man vs. machine may be a false dichotomy. Rather, it may be man and machine. ■

USING SOCIAL MEDIA TO FILL THE GAPS IN OBSERVATIONS DURING EMERGENCIES

Rebecca Goolsby, Ph.D., ONR Program Officer, Warfighter Performance Department

Guido Cervone, Ph.D., Department of Geography and Institute for CyberScience, The Pennsylvania State University; Research Application Laboratory, National Center for Atmospheric Research.

Novel hazards and disasters, especially complex environmental threats, are increasingly entering the Navy's sphere of operational responsibility. The U.S. Navy is increasingly called upon to providing advice on hazards and threats in the Arctic and the rising sea levels that threatened island communities, to study and maintain situation awareness of emerging health threats, and to provide disaster response and threat mitigation in a variety of novel situations. Combatant commands are beginning to prioritize situation awareness involving threats and hazards to human security from a disaster response operational perspective.

Hazards fall into three, often overlapping, categories: (1) natural hazards: events that naturally occur in nature; (2) man-made (sometimes called anthropogenic) hazards: events that are result from human activities or accidents; and (3) technological hazards such as the catastrophic collapse of infrastructure needed for society to function, such as roads, communication networks and power grids.

A major catastrophic event can have serious regional, even global impact on human security. It may claim thousands of lives and cause billions of dollars of damage, and create follow-on hazards that can render large territories uninhabitable, such as earthquakes which create tsunamis or epidemics that fuel food shortages. Many times, major catastrophes trigger substantive civil unrest and throw political systems into disarray. Further, the increasing urbanization of human society, including the emergence of megacities, has led to highly interdependent and vulnerable social infrastructures that may lack the resilience of more agrarian, traditional society. The proliferation of nuclear power facilities and waste storage, high dams, and other complex technological infrastructures that support dense population pose additional risks as they hold together complex systems needed to support urban societies. The ongoing Fukushima reactor crisis offers an important case in point, where a natural disaster engendered a technological one.

At the same time, novel information streams are redefining situation awareness. These streams are diverse, complex and overwhelming in volume, velocity and in the variety of viewpoints they offer. Never in history have people been able to know so much about our planet. Negotiating these

overwhelming streams is beyond the capacity of human analysts. Current research offers some novel capabilities to utilize these streams in new, groundbreaking ways, leveraging, fusing and filtering this new generation of air, space and ground-based sensor-generated data.

OBSERVATIONS AND THEIR LIMITATIONS

In recent years, advances in science have enabled new capabilities to observe the Earth and its environment through the use of air, space and ground based sensors. This has led to the generation of large, dynamic, and geographically distributed spatiotemporal data. Social media is a significant new information stream that can complete the picture by providing a rapidly updating dataset that not only complements the remote sensing observation, but also add an additional subjective view of how people perceive and react to hazards.

FUSING REMOTE SENSING, NUMERICAL MODELS AND SOCIAL MEDIA DURING EMERGENCIES

Remote sensing has become the de-facto standard in observing the Earth. Numerous air-borne and space-borne sensors provide an unprecedented access to high resolution spatio-temporal information about the Earth and its climate. During Humanitarian Assistance and Disaster Relief (HA/DR) Operations, these data are fundamental to understand the extent of the affected areas. Priority issues of quantification of damage, relief coordination and recovery operations also require this sort of dynamic data.

Despite the wide availability of large datasets from numerous sensors, specific data might not be collected in the time and space most urgently required. This may be due to satellite revisit time limitations, atmospheric opacity, or other obstructions. Geo-temporal gaps result,

for example, in the case of significant storms or remote locations. Complex locations often occur where sensors are not available or not tasked with sufficient frequency to capture the desired data for the necessary length of time. Social media can be used to 'fill the gaps' in these spatio-temporal data, and augment the initial satellite observations with tweets, photos or videos about an occurring event.

Harvesting and analyzing social media is a difficult and challenging task. It requires a truly inter-disciplinary team. In order to efficiently and effectively use these non-authoritative sources to augment satellite observations. Artificial intelligent and automated processes can provide many of these capabilities: accounting for missing data, filtering data for content, providing geolocation and other text mining and network analytic processes.

The data fusion problem is complex because of the extreme heterogeneity of the data: Remote sensing observations which have a high spatial but low temporal resolution, and social media which have a high temporal but low spatial resolution.

ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING

The rate at which geospatial data are being generated exceeds our ability to organize and analyze them. These developments are quickly leading towards a data-rich but knowledge-poor environment.

New challenges arise from an unprecedented access to massive amounts of data.

Geoinformatics algorithms are needed to address these scientific and computational challenges and provide innovative and effective solutions to analyze these large, often multi-modal, spatiotemporal datasets. Traditional data mining techniques do not incorporate the idiosyncrasies of the spatial

domain, which include spatial autocorrelation, spatial context, and spatial constraints. New methodologies are being developed for the fusion of geo-spatial data from remote sensing instruments with data from social networks. Artificial intelligence and machine learning are used to extract patterns from massive amount of data. The patterns are used to identify anomalies, cluster behaviors and to predict future outcomes.

2011 NUCLEAR ACCIDENT IN FUKUSHIMA, JAPAN

On 11 March 2011 at 05:46 UTC (14:46 local time, UTC +9) a massive Mw 9.0 underwater earthquake occurred 70 km offshore of the eastern coast of Japan, with epicenter at 38.322N and 142.369E. The earthquake generated a tsunami that rapidly hit the eastern coast of Japan, and propagated across the Pacific ocean to the western coast of the Americas. The tsunami wave hit the Fukushima power plant

Several radioactive releases ensued as a result of an increase of pressure and temperature in the nuclear reactor buildings. Some releases were the result of both controlled and uncontrolled venting, while others were the result of the explosions that compromised the containment structures. The explosions were most likely caused by ignited hydrogen, generated by zirconium-water reaction occurring after the reactor core damage.

The largest radioactive leaks occurred between

the 12 and 21 March 2011. Radioactivity was recorded at different locations throughout Japan on the ground, in the water and in the air. The individual radionuclide distributions measured over central-east Japan from the Fukushima nuclear accident indicate that the prefectures of Fukushima, Ibaraki, Tochigi, Saitama, and Chiba and the city of Tokyo were contaminated by doses of radiations, and that a large amount of radioactivity was discharged on March 15th and 21st.

The radioactive cloud was quickly transported around the world, reaching within a few days North America and Europe. Radioactive concentrations were also recorded along the US West Coast. Estimating the fate of the contaminants and predicting their health impact quickly became an issue of great importance.

Transport and dispersion (T&D) models can be used to compute radioactivity levels and ground deposition. Various models are available depending on the scale of

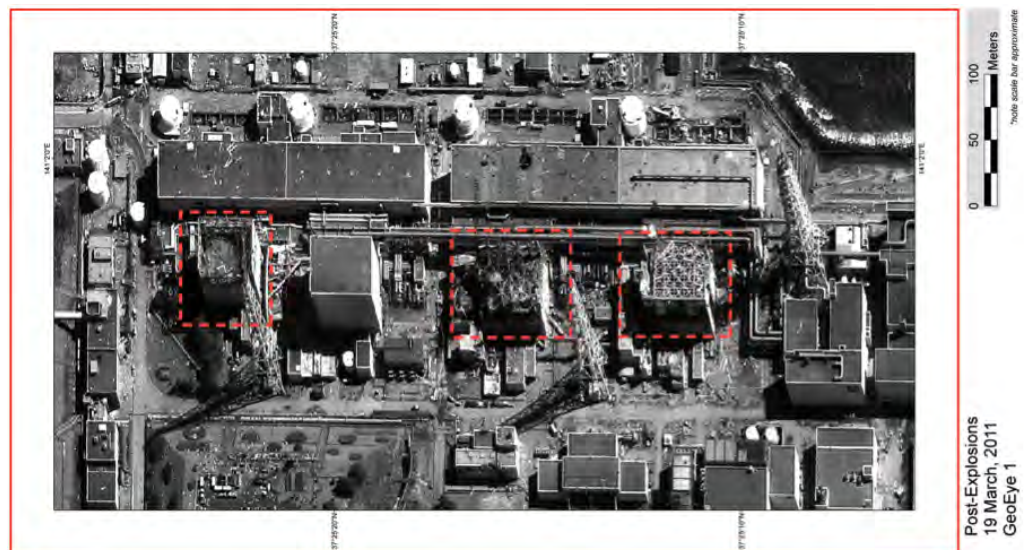


Figure 1.
Reactors damage at the Fukushima nuclear power plant observed by the GeoEye1 satellite on 19 March 2011 about 40 minutes after the earthquake, leading to the catastrophic failure of the cooling system.

the problem. For instance, the exclusion zone and its evacuation are determined by dispersion simulations in the area immediately surrounding the nuclear reactor on a scale ranging from meters to a few kilometers. Contamination at a planetary scale can be assessed by long range transport models.

with $n = 1 \dots N$. Figure 2a shows a sample steady plume represented as a sequence of identical releases with the same rate q (the area of each release is constant). A time-varying release rate Q_n (Fig. 2b) can be obtained by multiplying q by a scalar w_n for each release n . The goal of the reconstructing procedure is to determine the vector $W = \{w_1 \dots w_N\}$. The vector W is identified through a stochastic optimization process driven by an evolutionary algorithm that minimizes the error between the radioactivity levels measured at different locations in the domain, and the radioactivity simulated using the Lagrangian SCIPUFF transport and dispersion model.

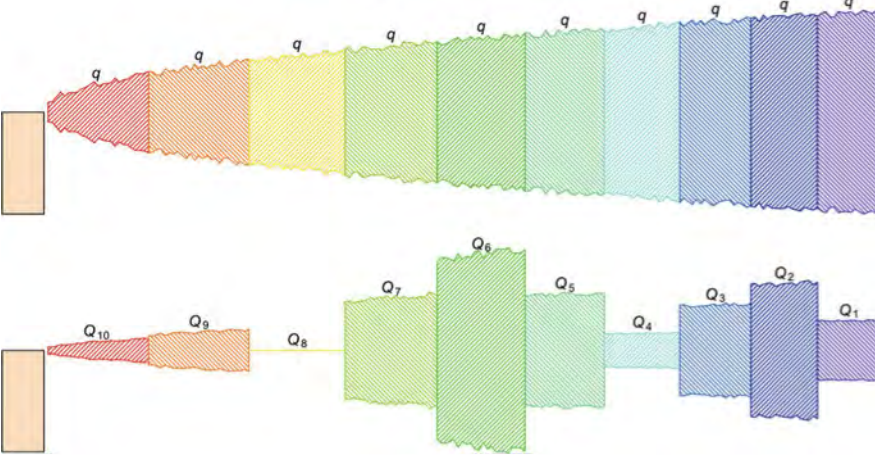


Figure 2.
Multiple consecutive releases of finite duration: releases of same rate q (top); discretized time-varying releases of different rate Q_n (bottom)

Dispersion simulations require meteorological data, terrain characteristics, source location and release rate. A major problem with the simulation of the Fukushima accident is the large uncertainty associated with the time-dependent release rate of radioactive contaminant.

ESTIMATING THE TIME-DEPENDENT RELEASE RATE

The methodology used identifies the source release rate which minimizes the error between observed radioactive concentrations and those resulting from a simulation process. In order to be able to reconstruct an unsteady release rate, a continuous release with a virtual constant rate q is discretized into a sequence of N consecutive finite-duration releases Q_n ,

THE IMPORTANCE OF SOCIAL MEDIA

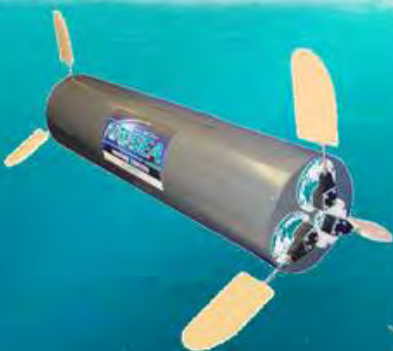
The time-dependent radioactivity (or concentration of radioactive material) is thus simulated at each ground location where radioactivity measurements are available. Inherently, the more ground concentrations available, the more accurate the solution. Unfortunately, the data available contains both spatial and temporal gaps due to missing or incomplete data, limitations of the sensors or of the platforms used for the collection.

Filling these gaps is necessary for the methodology to work and give accurate results. A very dense network of privately owned radioactivity measurements was established in the immediate aftermath of the explosion. Data from these sensors can be used to better estimate the non-steady release rate.

Furthermore, it is important to understand how people react to such a catastrophic and potentially life threatening event. Monitoring social media, such as Twitter, it is possible to understand the dynamic evolution of how citizens perceive an event, and react to potential danger. The combination of

accurate simulations for the contamination of plume, paired with an understanding of people reaction to the developing event can be used to optimize HA/DR operation, protect people, properties and the environment. ■

Swim Like Fishes: **BIO-INSPIRED UNDERWATER VEHICLES**



BIO-INSPIRED AUTONOMOUS SYSTEMS

Neural-based intelligent systems embedded into bio-inspired robots and vehicles with superior mobility and the ability to operate autonomously in complex environments.

Thomas McKenna, Ph.D., ONR Program Officer, Warfighter Performance Department

Harbors and inland waterways can appear to be placid spaces, with the silence occasionally punctuated by the noise of powerboats and jet skis that echo off moored warships or cargo ships with unknown cargoes. Underwater, marine life moves silently, effortlessly wending their way among pilings and debris. Fish hover and glide, using their flexible fins to sense obstacles in murky waters. There is not a drop of water in the harbor, not an inch of bottom or breakwater or

pier that living organisms do not visit. Under the water there are many niches where a wide range of animal propulsion systems with special bio-senses are constantly on patrol, agile and silent. These animals sense currents, eddies, electric fields, subtle mechanical events, chemical gradients, polarized light and the motion of bodies both large and small.

Soon, a new generation of bio-inspired underwater vehicles may be silently patrolling harbors and rivers. These new autonomous vehicles will stealthily examine each ship and dock, perhaps leaving markers, and standing guard against divers that could pose a threat. Smaller bio-inspired vehicles will closely inspect complex ship structures like propellers and swim in tanks inside of ships. These new vehicles will have propulsion and control that make them stealthy, efficient for endurance and highly agile. ONR's Warfighting Performance Department, supported researchers seek to identify the principles, strategies and mechanisms used by aquatic animals and exploit these to achieve capabilities beyond the current engineering of undersea vehicles. For example, marine animals have a broad range of propulsion modes and ONR has been actively exploring many of them. Fish that swim at high speeds typically use the caudal part of the body for propulsion. This type of propulsion, called lunate tail propulsion, is dominant in tuna, which can achieve speeds of 50 knots, and we have been experimenting with robotic tuna since the 1990's. The latest version, called Ghostswimmer, is a complete platform already in use for ship hull inspection. An undulating propulsion like this is driven by an oscillating central pattern generator, like that seen in the nervous system in all vertebrate animals. However, to maneuver the Ghostswimmer an artificial neural network has to be trained over a range of swimming behaviors. Fish like perch, that are highly maneuverable and can hover and move backward, commonly have paired oscillating

fins in addition to the caudal fin. When the paired, or pectoral fins, have rays with a membrane stretched between them, as in perch fish, the fins make a complex motion on each beat that we are still exploring in the lab.

Many sea animals have solid pectoral fins, like sea turtles and penguins, and the hydrodynamics of foils that move like these fins have been well studied in the Navy lab at Naval Undersea Warfare Center (NUWC) Newport. Pectoral fins like this have both a pitching and heaving motion shifted in phase and these fins function as high-lift foils. High-lift foils produce substantially greater lift than rigid foils with a fixed angle of attack, such as traditional propellers, and they also produce substantially less noise. But bio-inspired high-lift foils have many degrees of freedom and a vehicle might have 4 or more foils. This produces a challenge for traditional engineering control approaches. Fortunately, nature provides a solution to this control problem. One of the important brain structures of the mammalian motor control system is the cerebellum which has a companion structure, the inferior olivary nucleus. Neuroscience and nonlinear dynamics analysis of this system show that it can generate a complex sequence of motor commands to multiple actuators. When a model of this system was built as neuromorphic electronic circuits it was found to be a remarkably effective controller for the vehicles that have been built using the high-lift foils. The highly maneuverable and efficient NUWC RAZOR vehicle (a multi-mission expeditionary autonomous unmanned underwater vehicle) has 4 of these foils.

There are other modes of animal propulsion worth exploiting. The Ghostfish lives in murky Amazon eddies and hunts completely by electrosense. This fish has a single long fin that extends like a ribbon along its ventral surface and it is propelled forward and backward by waves of undulations along this fin. ONR

supported researchers have bio-engineered the bio-inspired active electrosense to a practical device, and have built prototypes of this propulsion mechanism. These small vehicles might some day serve to inspect hulls and tanks for corrosion by sensing electric currents and conductivity changes.

Batoid fishes (e.g., rays) utilize one of two modes of locomotion, employing either undulatory or oscillatory kinematics along the body edge. ONR's Sea Warfare and Weapons Department supported researchers have explored the biomechanics in detail and built promising lab prototypes of a manta ray inspired vehicle.

Lastly, fluid jetting provides impulsive locomotion for animals like squid and efficient locomotion for jellyfish. University researchers are working hard to mimic these organisms.

In aquatic animals, the locomotion and sensing systems are complementary. We

"Don't be too surprised if someday a vehicle with a bat-like brain is flying underwater."

are just beginning to understand how the special senses like bio-sonar, lateral line mechanosensory and electrosense are integrated with the motor control in these animals. In aquatic animals, experimentally determining higher sensory representations (i.e., how these animals map the world,) is exceedingly difficult. But we can apply the lessons from other animals that live in a 3D world, like bats. Bats have a highly developed biosonar for hunting and navigation, which ONR supported researchers have been studying since the earliest days of ONR. More recently, researchers have been able to record from the individual neurons in the brains of

bats while they are flying through obstacles, and we are beginning to understand how their brains represent 3D space. Don't be too surprised if someday a vehicle with a bat-like brain is flying underwater.

Perhaps we will see bio-inspired autonomous vehicles inhabiting many aquatic niches, like their animal models. Several of the vehicles described above are already in demonstration phases, and ONR research programs are providing the Navy with important new options for future undersea vehicles. ■

ONR STEM Grand Challenge:

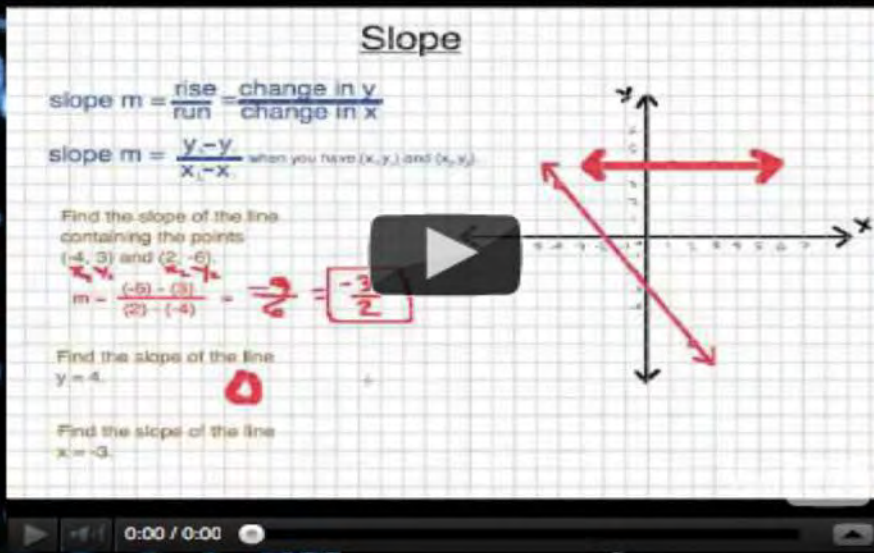
**ADAPTIVE,
GENERALIZABLE,
INTELLIGENT TUTORS**

**FOR STEM AND NAVAL TRAINING
AND EDUCATION**

Ray Perez, Ph.D., ONR Program
Officer, Intelligent Tutor Systems,
Warfighter Performance Department

LCDR Brent Olde, Ph.D., ONR Division
Deputy, Human & Bio-Engineered Systems,
Warfighter Performance Department

Only 16 percent of American high school seniors are proficient in mathematics and interested in a Science, Technology, Engineering, and Mathematics (STEM) career. Even among those who do pursue a college major in the STEM fields, only about half choose to work in a STEM related career. The United States is falling behind internationally, ranking 25th in mathematics and 17th in science among industrialized nations. Today's approaches to education and training must



change if our Nation is to be competitive in the dynamic global workforce.

“.. Leadership tomorrow depends on how we educate our students today—especially in science, technology, engineering and math.”

PRESIDENT BARACK OBAMA, SEPTEMBER 16, 2010

Due to the above data, Navy leadership has issued new policy and guidance focused on STEM. Based on a White House short list of STEM priorities, the 2011 Chief of Naval Operations' Guidance (CNOG 2011) specifically called on the Navy to focus on advancing STEM education in order to provide the technical skills required to lead tomorrow's Navy. To attain this goal, our recruits must have basic math and engineering knowledge,

technical acumen, as well as the ability to reason and make informed decisions using scientific knowledge, methods, and principles gained through a robust STEM based curriculum.

Today's teaching methods still rely on a typical schoolhouse approach, with one instructor teaching many students. This approach has severe limitations, it tends to leave advanced students bored and under performing students left behind and disengaged. Numerous studies have shown that a more effective way to teach all students is through one-to-one interactions between teachers and students. Such teaching methods lead to two standard deviations improvements in performance – roughly two letter grades – compared to the traditional classroom approach. However, it is difficult and costly to provide this level of individualized interaction to students, especially to students in overcrowded and understaffed schools.

Current advances in technology can help provide this individualized education but for it to be truly effective, it will require schools and teachers to change the way they do business. Some schools have recognized the need to reform the current approach to education and are experimenting with these educational technologies to augment classroom instruction throughout the United States. These technologies include “intelligent” tutors, video-games, simulations, computer-based testing, online education, and compelling forms of immersive learning environments. Technology-based learning tools allow students the flexibility and convenience of education and training by tailoring the pace, content, and course sequencing appropriate for each individual. Moreover, numerous studies document that online or mobile computer-based, simulation-based, game-based, and intelligent tutoring technologies all positively improve student learning effectiveness.

One of the most consistent findings comparing technology-based instruction with conventional instruction is that technology-based instruction affords significant savings in time spent learning the material. As early as 1977, Orlansky and String found a 54 percent average reduction in time to reach instructional objectives in their review of 13 technology-based military training programs. Fletcher (2001) reported an average time reduction of 31 percent in 6 assessments of interactive multimedia instruction applied in higher education, and Kulik and Fletcher (2013) and Kulik (1994) reported time reductions of 34 percent in 17 assessments of technology used in higher education and 24 percent in 15 assessments of adult education.

Overall, it seems reasonable to expect technology-based instruction to reduce the time it takes students to reach their educational objectives by about 30 percent.

However, we have gone a step further, going beyond interactive multimedia and into the realm

of Intelligent tutors. While specific Intelligent tutoring systems vary, these software programs (grounded in our latest knowledge of how learning occurs) mimic the actions of human tutors. These “tutors” enable students to work through problems in their own way and provide prompts, hints, and feedback as needed at every step of a student’s progress (Kulik & Fletcher, 2013).

To address the current gaps in technical tutor capabilities, the Office of Naval Research (ONR) launched the 2013 STEM Grand Challenge, *“ADAPTIVE, GENERALIZABLE INTELLIGENT TUTORS FOR STEM AND NAVAL TRAINING AND EDUCATION.”*

The objective of the STEM Grand Challenge is to develop cost effective, generalizable instructional tutors for STEM training and education that will raise a student’s performance by at least two letter grades. Achieving this objective requires a concerted, multi-disciplinary effort to place science based instructional technologies into the hands of students across the socio-economic spectrum and may encompass a broad range of capabilities including desktop, mobile, and gaming platforms. The STEM Grand Challenge provides the right blend of incentive, risk and benefit to revolutionize the state of the art of intelligent tutors. Its goal is to develop innovative, scalable, and affordable technologies that blend the best teaching approaches with cost effective design solutions capable of broad dissemination to a wide range of students.

The Grand Challenge is a three year, two phase effort. Phase 1 will develop an intelligent tutor technology for a STEM curriculum, targeted at the middle-high school level. The team(s) that demonstrate cost effective, student improvements will progress to Phase 2. In this phase, the tutor will be modified and implemented to address specific Department of the Navy training and education needs.

RESEARCH TEAMS FOR PHASE 1 ARE:

- University of Memphis, Dr. Xiangen Hu and Dr. Art Graesser: Pre-algebra and algebra.
- Arizona State University, Dr. Kurt VanLehn / University of California Los Angeles, Dr. Eva Baker: Domain model construction
- University of Massachusetts, Dr. Beverly Woolf / Worcester Polytechnic Institute, Massachusetts, Dr. Neil Heffernan: Middle and high school math.
- Raytheon BBN Technologies, Dr. Bruce Roberts and Dr. Rohit Kumar / University of California Los Angeles, Dr. Eva Baker: Advanced Physics.

The STEM grand Challenge contracts were awarded January 2013, all four contractors have a history of developing effective tutoring systems. Their new efforts will focus on improving the effectiveness of their current tutors by utilizing new innovative technologies and science of learning concepts to enhance the tutor's pedagogical approach. Their secondary task of developing authoring aids will reduce the time and cost of producing new intelligent tutors thus stimulating early transition. The primary objective of this effort is to build efficient and effective intelligent tutors to teach highly complex STEM knowledge and relevant skills sets to students at all ability levels. Seen more globally, these technologies will produce a workforce who have basic math and engineering knowledge, technical acumen, and the ability to reason and make informed decisions using scientific knowledge, methods, and principles.

For the Navy, the target audience for these tutors are (a) incoming Navy recruits that lack foundational knowledge and skills in STEM, (b) returning veterans entering intuitions of higher education, vocational training, and (c) middle and high school students who will make up the majority of tomorrow's workforce.

IMPACT

As Navy equipment and technical jobs continue to become increasingly more complex the Naval workforce will require increased knowledge and skills in STEM subjects to meet the technical demands of current and future jobs. Modern jobs require continual education, so learning needs to be faster and cheaper. Job relevant knowledge and skills are dynamic and need to be durable throughout their career. So education and training and retraining programs will need to be available anytime, anyplace, and anywhere on demand. Intelligent Tutoring Systems will be able to provide instructional programs that are easily updated and deployable while at the same time reducing the time it takes students to reach their educational objectives by at least 30 percent.

CLASSROOM OF THE FUTURE: A VISION

"The classroom of the future" most certainly will be loaded with technology from the use of 3D technology, natural language interactions, and your own personalized avatar. Imagine a biology student looking at a heart on a 3D screen or even printing one out to study the underlying mechanisms. Internet access could be a full sensory experience with the use of "Google" Glasses, the Apple i watch, Bluetooth earpiece and voice recognition system. Imagine a star trek-style halo deck where students could experience history in 4D. However, the most dramatic change is that education and training will be the role of the teacher/instructor evolving from "sage on the stage" to that of a coach, augmented with ITS in every subject. Education and training will become personalized where students will have their own personal learning assistance (PAL) that

incorporates intelligent tutoring technologies with the capability to record courses taken, grades earned, and recommend future courses required to reach personalized goals. Just imagine your PAL not only tracks your learning progress but effectively adapts instruction to match your particular learning style/preference. This PAL can handle basic information delivery and correct common misconceptions, providing time back to instructors/teachers who can provide targeted help with their limited time. ■

Figure 2.

The WalkPoly problem is part of Dr. Beverly Woolf's MathBuds Intelligent Tutor. It presents mathematics and offers additional help if the student experiences difficulty. This screen shot shows a student working to calculate the parameter of a parallelogram and about to click on the "Hint" link for assistance.

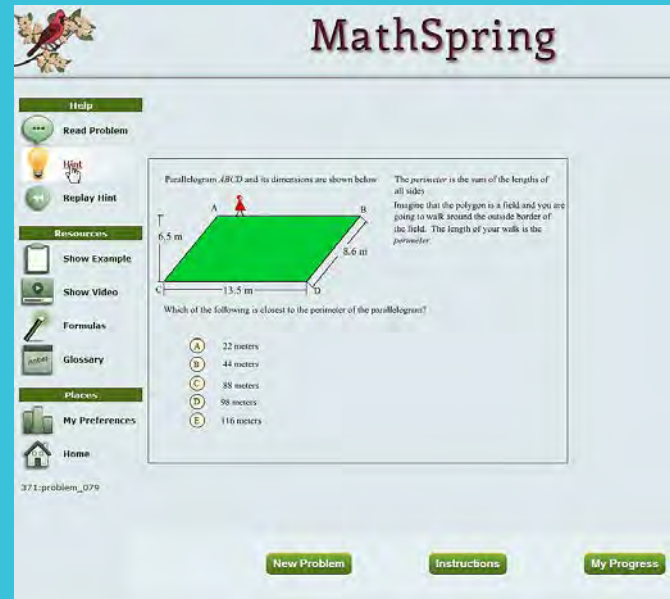
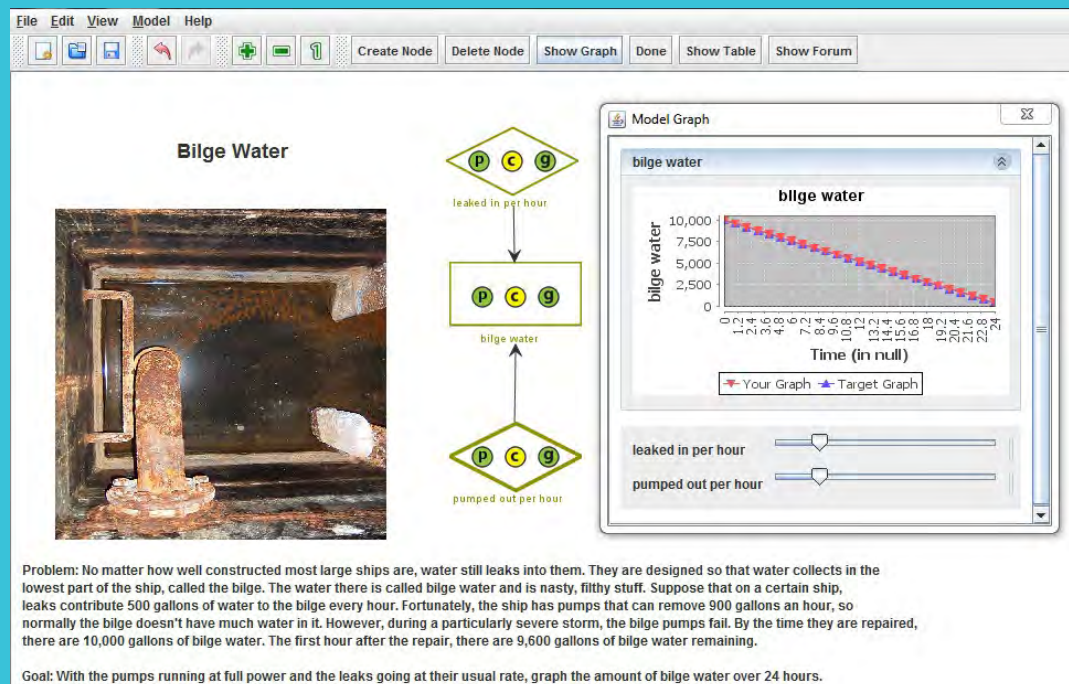


Figure 3.

The Bilge water problem is part of Dr. Kurt VanLehn's Intelligent Tutor. The image displays a model where the student can adjust parameters of the leak and the ship's pump capacity. Using this model the students can explore key parameters of the model to see how they relate.



APPLYING SUBMARINE TACTICAL CENTER KNOWLEDGE TO

UUV AUTONOMY

ATLANTIC OCEAN (June 30, 2011) The Virginia-class attack submarine USS California (SSN 781) underway during sea trials. (U.S. Navy photo by Chris Oxley/Released)

Mark Rothgeb, Applied Research
Laboratory, The Pennsylvania State University

John Sustersic, Ph.D., Applied Research
Laboratory, The Pennsylvania State University

Autonomously operating Unmanned Underwater Vehicles (UUVs) are beginning to receive a larger focus in the Navy due to continuing technological advances and the achievements of unmanned aerial vehicles (UAVs) like the Predator. The Predator's huge impact in the Afghanistan and Pakistan operational theatres is largely due to the ability to achieve safe standoff for the warfighters and for manned and unmanned aircraft to share the sky while reliably performing precision strike operations.

Today's autonomous vehicles are capable of extending the reach of the Navy but still depend largely on remote operation by humans either in-the-loop or on-the-loop. Thus, the term "unmanned system" is today a misnomer in the sense that there are substantial human resources required in their operation, particularly to provide higher-level cognitive reasoning and decision-making.

The challenges of the maritime undersea environment are unique and particularly severe as noted in the Defense Sciences Board's task force report:

"...technology cannot overcome certain physical limitations of the marine environment, essentially mandating greater autonomy...As improvements are made in energy density/endurance, unmanned maritime vehicles will be able to conduct far-forward missions, both enabling and capitalizing on future advances in autonomy."

Though advances are currently enhancing autonomy in real-world problem domains, the wealth of knowledge that resides within the submarine community has not been tapped. With over 100 years of experience dealing with the manifold unpredictable realities of the ocean environment and underwater operations, the submarine community has Subject Matter Experts (SMEs) who have served for years in different roles, passing their experience on through rigorous training, mentoring, and in-situ operational oversight.

(1) Knowledge Capture: A broad spectrum of active and retired Commanding Officers (COs), Executive Officers (XOs), Officers of the Deck (OODs) were interviewed including Navigators and Department Heads. The interviews included 10 individuals who served on a total of 26 submarines with over 150 years total experience. Figure 1 depicts the focus of this task which resulted in a knowledge base capture of their relevant experiences.

The interviews focused on common operations (e.g., ocean transit and coming to the surface). Decision-making processes were captured along with how the SMEs “built a picture” of the situation and what they did to sufficiently fill in knowledge gaps.

Preliminary analysis suggests that the knowledge across individual COs can be aggregated into an artificial “mission commander” that could be represented in a cognitive architecture such as ACT-R (Adaptive Control of Thought – Rational) or Soar (State, Operator And Result), adapted for an AUVs capabilities. Though an AUV will operate differently than a submarine, some knowledge will be directly relevant (e.g., waterspace management), other knowledge may be adapted for an AUV (e.g., coming to the surface), and some will not be relevant (e.g., maintaining crew alertness). AUVs will perform limited tasks that it can do reliably while reducing warfighter risk.

(2) Autonomy Architecture: Using the submarine tactical center as a model for autonomy is a potentially groundbreaking change for future AUV systems with significant potential benefits. Traditionally, a tradeoff analysis is performed early in the design of a new system to select a single autonomy solution for the entire system which may be (a) a reactive, behavior-based system as represented by the Prototype Intelligent Controller (PIC), (b) a hierarchical or deliberative system as represented by the Maritime Open Autonomy Architecture (MOAA) and the Control Architecture for Robotic Agent Command and Sensing (CARACaS), or (c) a cognitive system as represented by Soar and ACT-R. The difficulty is that certain technologies apply better to certain

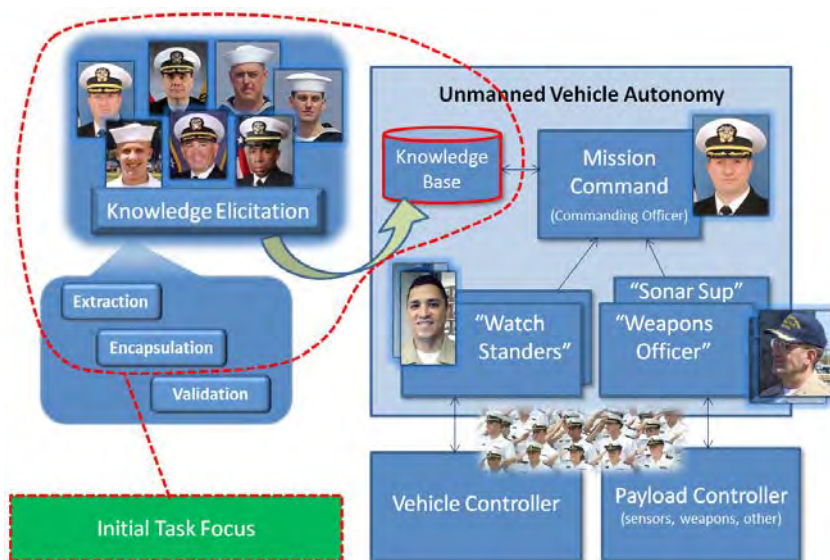


Figure 1.
Knowledge Elicitation

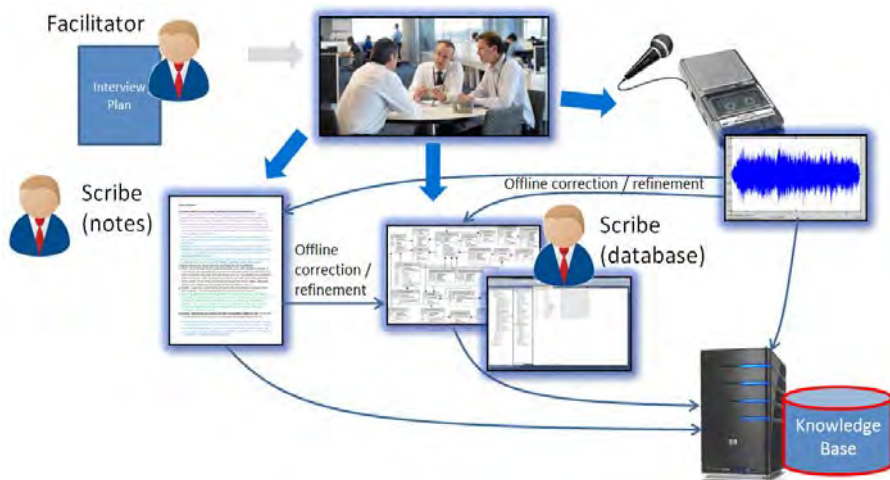


Figure 2.
Interview Process

tasks. A submarine navigator builds a plan in a deliberative manner taking in knowledge such as past plans, mission context, current environmental conditions, and geographical features. Collision avoidance while in transit is more suited to reactive systems. Furthermore, the CO operates in a cognitive decision-making role. The proposed architecture allows the most relevant technology solution for each role.

Furthermore, the model naturally modularizes the autonomy architecture into distinct functional roles that map directly to existing

operational roles and responsibilities that are well-defined and understood by the submarine community. The interactions (messages) between the watchstanders is known and precisely defined in the Submarine Interior Communications Manual (ICM). This intrinsic modularity enables interchangeability and re-use – e.g., a different “Navigator” algorithm can be replaced without impacting the rest of the system. Functional decomposition enables a mission specific “CO” module that leverages the standard piloting, navigation, vehicle control and payload operations. Standards-based interoperability is achieved by message-based interactions of the various roles based on the ICM. Thus the autonomy architecture is understandable and accessible to operators and maintainers - unlike many existing autonomy systems.

The resulting autonomy architecture breaks the paradigm currently used today in every AUV system and has the potential to position the Navy with a robust architecture for future AUV systems.

(3) Cognitive Systems: Today’s state-of-the-art in cognitive software architectures is inadequate for the reliable operation of truly autonomous systems, especially in dynamic and often unconstrained environments typical of the modern battlespace. The challenge faced today centers on two key aspects: how to express operationally complex mission decision-making processes in some form that can be executed by an autonomous system and how to transform this expression into executable code.

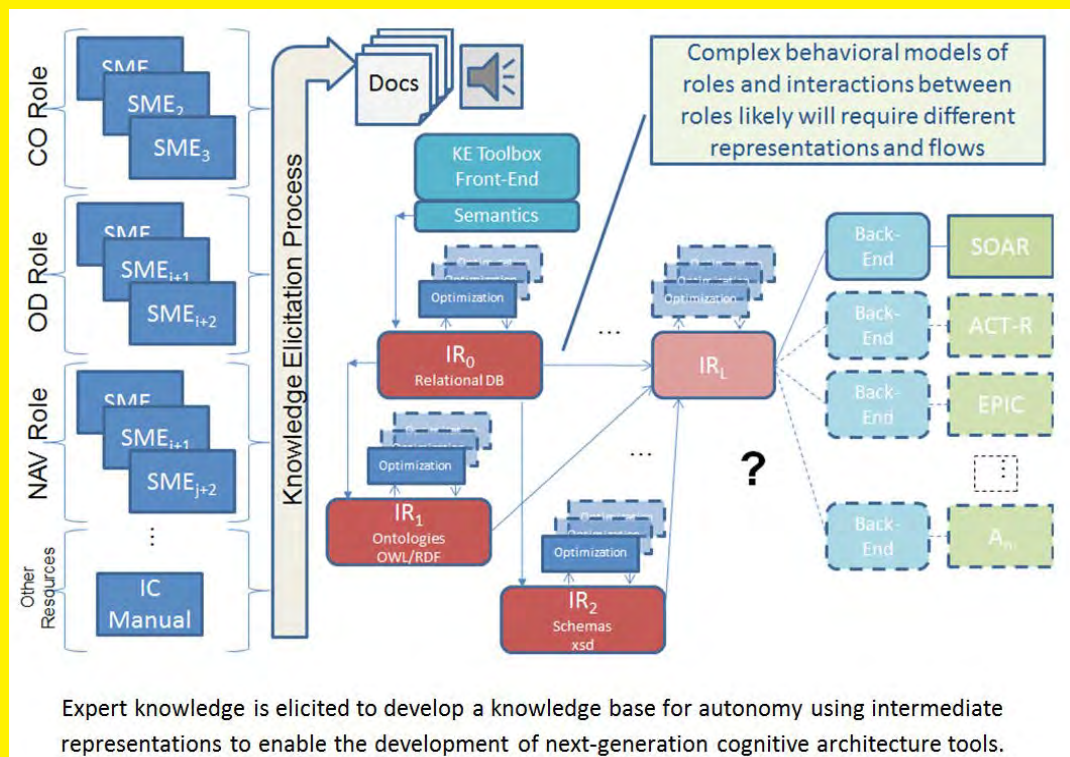
These challenges are analogous to the problem solved in many modern programming languages that leverage the capabilities of different hardware architectures while avoiding the impossible costs of providing a compiler for each language-architecture pair. Decades of research has led to a sustainable and cost-effective solution – utilizing a common Intermediate Representation (IR) that requires only a customized ‘front-end’ for each language and a customized ‘back-end’ for each target architecture, enabling today’s excellent programming tools and the integration of codebases written in different programming languages into complex applications.

Today's challenge in autonomy is parallel: (a) there are many cognitive architectures (e.g., SOAR, ACT-R, neural networks, etc.) and differing levels of expression (symbolic vs. connectionist), each with differing strengths and weaknesses; (b) programming any of these is painstaking, time-consuming, and requires expert programmers who often don't fully understand the application domain; (c) there are many autonomy platforms, each with differing hardware architectures; and (d) there are few tools available to assist in programming, debugging, and maintaining these cognitive architectures.

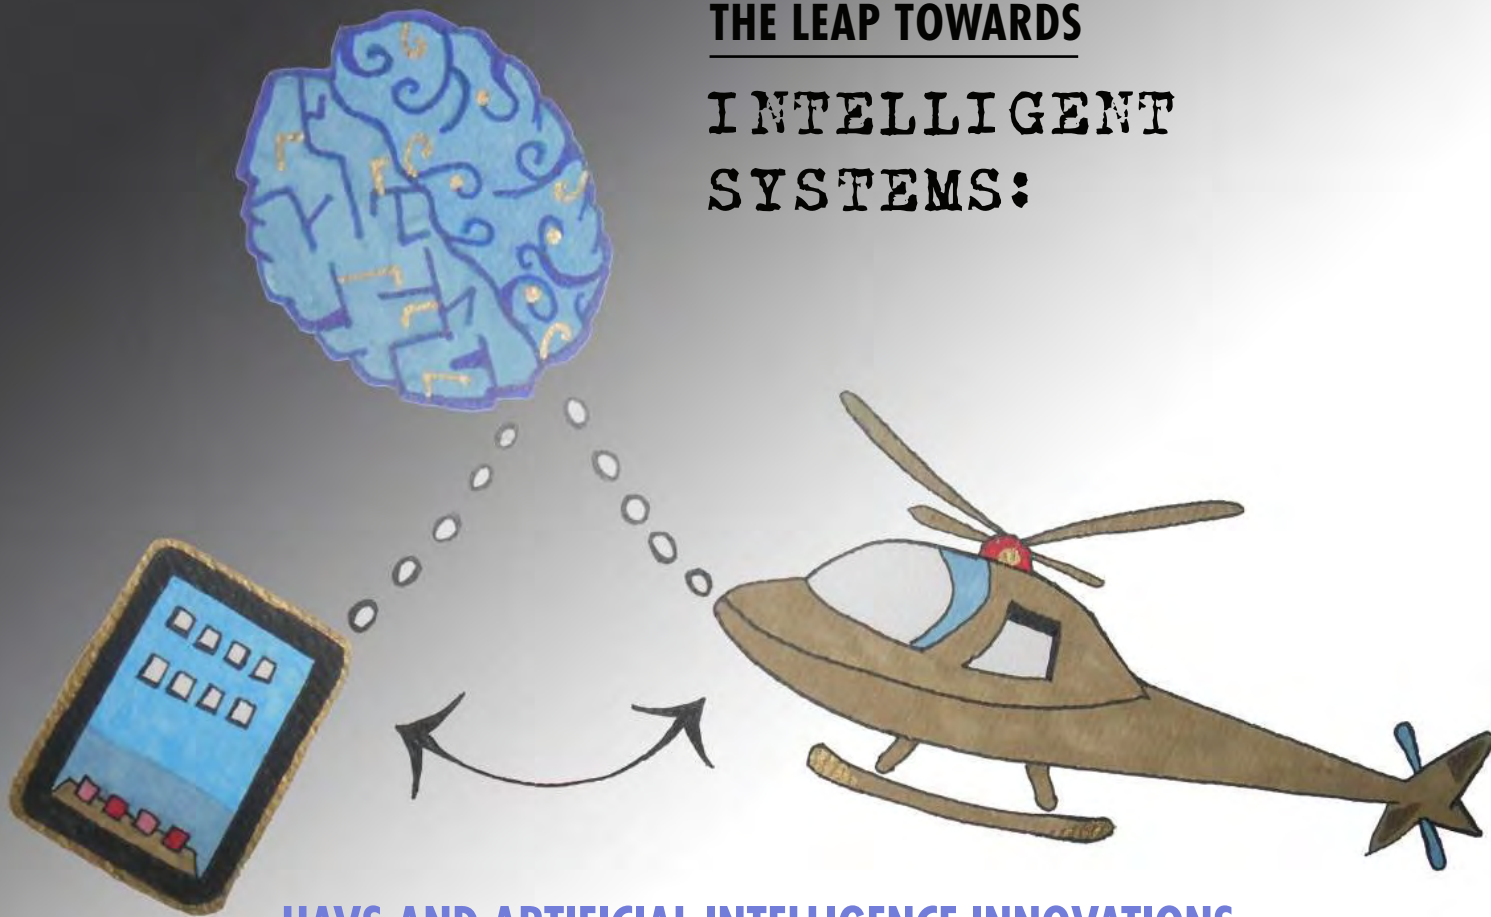
The lessons learned in compiler development can be applied to the challenges of advancing the state-of-the-art in autonomy by taking a similar approach in which different levels of IR are used to allow greater flexibility and cost effectiveness. Figure 3 shows this compiler concept as applied to cognitive autonomy. The challenges here are even

greater in that the knowledge domain and language used to express human subject matter expertise is not easily represented formally. Moreover, different aspects of knowledge map to different kinds of representations. For example, messaging is well-structured for representation as XML Schemas, while the behavioral modeling, decision making processes, etc., are almost certainly expressed in other methods. This approach enables system designers to efficiently employ different types of interoperating cognitive architectures in various roles according to the operational requirements of the role. Ultimately, this approach could lead to a Cognitive Autonomy 'debugger' capable of tracing back through the various representations to find the source of errors, allow better analyses and the development of end-to-end tools to support cognitive architecture development, and bring programming and maintenance of complex cognitive architectures into the reach of mainstream programmers and eventually operators in the fleet. ■

Figure 3: The Cognitive Autonomy Compiler Concept



THE LEAP TOWARDS **INTELLIGENT SYSTEMS:**



UAVS AND ARTIFICIAL INTELLIGENCE INNOVATIONS

Azi Sharif, Ph.D. Candidate, ONR Program Officer,
Naval Air Warfare and Weapons Department

An unmanned aerial vehicle (UAV), is known as an aircraft that can fly without a human pilot on board. It is traditionally controlled by trained human operators at the Ground Control Station. The flight path is controlled remotely using way point navigation. Traditional UAV's are not necessarily smart, and do not have the ability to analyze, reason, and make their own decisions; they do as they are told. Therefore, traditional UAV's require intense training, operator involvement, and are not reliable to function for long hours or in dangerous dynamic conditions and harsh terrains.

The next generation of unmanned aerial vehicles requires significantly greater sensing, perception, reasoning, and decision making capabilities. They should have the ability to make their own decisions while supervised by a field operator on the ground. The data that the UAVs collect will need to be processed quickly on-board the aircraft without the need for humans to spend hours going through and sorting the data. The field operator will be capable of selecting missions, desired supplies to be delivered, and location, quickly and easily by using an iPad like device. The UAV will have the ability to deal with uncertainty. The so called "Brain" of the future UAV is enhanced and developed by algorithms and artificial intelligence.

WHAT IS ARTIFICIAL INTELLIGENCE (AI)

AI is “the study and design of intelligent agents.” (Norvig, 2003) The intelligent agent can be a simulated ant, a robot, aerial system, ground system, underwater system, or simply an App on your personal iPad. Artificial Intelligence enables systems to be capable of perceiving the environment, reason, make decisions, and take actions. Some of the enduring challenges of AI are also the capabilities that it enables including knowledge, reasoning, planning, learning, communication, and perception.

DISRUPTIVE TECHNOLOGIES: SOME OF THE LATEST SYSTEMS WITH ARTIFICIAL INTELLIGENCE

Researchers at Boise State University are building a Computer Chip Based on the Human Brain on a project funded by National Science Foundation. The team plans to design the next generation computing chips that mimic how the brain processes information and its pattern recognition capabilities rather than the traditional computer chips. This could potentially enable systems the ability to learn, adapt and respond to their environment.

Another great example of the use of AI is a new program which aims to develop a squad of unmanned underwater rovers deployed from a surface vehicle. While using artificial intelligence, the rovers would have the ability to cooperate and communicate among their group and cope with the uncertainties of the marine environment. This could be further enhanced by extension of the cooperative behavior beyond just the underwater rovers. The group could consist of underwater rovers, robotic aerial and ground systems all coordinating on a joint mission to collect information.

NEXT GENERATION LEAP TOWARDS AUTONOMY: THE AUTONOMOUS AERIAL CARGO UTILITY SYSTEM INNOVATIVE NAVAL PROTOTYPE (AACUS INP)

The Office of Naval Research in Washington, DC, launched a program in September of 2012 to develop autonomous capabilities for robotic helicopters. The technologies developed under AACUS will be a huge leap in autonomous aircraft capabilities. Two contracts were awarded to Lockheed Martin and Aurora Flight Sciences on Sept. 28, 2012 to develop robotic rotorcraft capable of supporting rapid autonomous aerial cargo delivery to the battlefield. In addition to industry partners, ONR has teamed with NASA’s Jet Propulsion Laboratory and the U.S. Army Aviation and Missile Research Development and Engineering Center. AACUS is a fast paced program developing a cross-platform software and sensor package focused on game changing intelligent autonomy, with a novel, iPad-like interface device.

The system will support Navy and Marine Corps units under hostile conditions and could be operated by any warfighter on the ground with a smartphone device. AACUS originated from Novel Sensor Suite

Two flight demonstrations will take place within 18 months of program initiation currently scheduled for Spring 2014. The helicopter will be flying autonomously with all data processed on-board the aircraft, selecting its own landing site, and landing while supervised by a human on the ground via an iPad-like device vice the traditional remote controlled or way point controlled UAV.

Although the scope of this program is focused on the mission of cargo resupply and delivery, after successful implementation of the innovative autonomous technologies, the future technological advances and possibilities are endless. For example, within the next 10 years, a smart robotic helicopter could have the capabilities to avoid being shot at, land in a battle zone, deploy a smart robotic ground vehicle to pick up an injured soldier and place the soldier on-

board the helicopter, and fly them back to base. The advances in healthcare and patient monitoring systems in the near future will also be among some of the most game changing developments. Cell phones will be able to scan patients and detect and monitor health problems. Robots will be able to track blood pressure, and smart pills and patches will be able to monitor patient's reaction to a drug. All this information will be transmitted to the doctors via an iPad-like device with a simple user interface. In addition, a network of smart helicopters, ground vehicles, and underwater vehicles, could work cooperatively to achieve their task and missions under dynamic and hostile conditions while supervised remotely by human operators. As we leap towards more intelligent and adaptive systems, the potential technological advances in artificial intelligence and unmanned systems becomes possible.¹ ■

Director's Corner:

THE EVER GAME-CHANGING FIELD OF **Artificial Intelligence**

Larry Schuette, Ph.D., ONR Acting Director of Research

Imagine if after every time your favorite football team scored a touchdown, the rulebook allowed us to move the goal post and the end zone back, making the possibility of a touchdown more unlikely and eventually impossible with every impending score. That foredoomed scenario is exactly what the research field of artificial intelligence (AI) has been battling ever since the first discoveries in the early 1950's. Many AI researchers have called their field of expertise a curse; the minute they are successful in their research, that goalpost is pushed back a little further and they are reminded that they have yet to deliver AI.

At this point in the newsletter, you're very familiar with AI and autonomy and the vast amount of applications these technologies have on our naval systems and platforms. To recap, AI is a cluster of scientific, engineering and philosophical endeavors guiding the development of "smart" machines. To truly create smart machines, we need to make significant advancements in the following areas: perception, learning, knowledge representation and reasoning, planning and acting, communication and interaction, and robotics. Examples of AI success stories in these capability areas are technologies that you are very familiar with like GPS, Siri and Google Voice, Roombas, TSA scanners, and UPS delivery. All of these examples are technologies that achieved significant breakthroughs in AI, however, once realized, demonstrated and commercially available, all of a sudden they are no longer recognized as AI accomplishments. Goal post pushed back once again.

Since early Greek mythology and perhaps before that even, thinking machines and artificial beings have been a part of our stories. There's Talos of Crete from Greek mythology, Mary Shelley's

¹ Cited work: Russell S. and Norvig P. (2003), Artificial Intelligence: A Modern Approach. Prentice Hall, New. York

Frankenstein, Starwars' R2-D2 and C-3PO, and Hal from 2001: A Space Odyssey. However, it wasn't until 1941 and the development of the first electronic computer that we had the technology to make fiction, non-fiction. Even with all of those stories, it was nearly a decade after the first electronic computer was invented that connections between human intelligence and machines were first made with the realization of feedback theory. There were significant advancements and investments over the next 20 years with the Department of Defense taking a strong interest in furthering this field. Unfortunately, and like so many fields of research, AI experienced a dry season in the 1970's where there were very few projects and limited funding due to the scale of some of the problems the researchers faced and the ever growing field of skeptics. After a revival in the 1980's and an even longer "winter" season, the field really took off in the 1990's due to Moore's law and a greater emphasis towards solving specific AI research sub problems.

In 1997, IBM put their AI research to the test with their computer, Deep Blue up against reigning world chess champion, Garry Kasparov. Since then, we've seen phenomenal accomplishments such as the DARPA Grand Challenge and the robot that drove 131 miles on a desert trail without human guidance or pre programming. We've witnessed that same concept further developed with a robot responding to the DARPA Urban Challenge by driving 55 miles in an urban environment with our daily traffic hazards like stop signs, traffic, pedestrians, etc. In 2011, IBM was back at it with their Watson system which defeated two of Jeopardy's all time winning champions. You'd be hard pressed to find a kid in American who hasn't heard of Microsoft's Kinect system for the Xbox 360 that emerged from lengthy AI research.

We all enjoy and benefit from the technology made possible due to substantial AI research, whether we know it or not. This is truly a field that has produced game-changing technologies – and promises to deliver more as we take the smart machines of today and make them even better, faster, and smarter – for tomorrow. ■

LAST ISSUE OF THE INNOVATION NEWSLETTER:

This volume of the Innovation Newsletter will be the final issue and marks an important reorganization at ONR. We are in the process of reorganizing the Director of Innovation, Research, and Transition into two Directorates: Research and Technology. I'm honored to take over as Director of Research, and I look forward to working closely with the Director of Technology, Dr. Tom Killion. We've continued to recognize the value that this publication provides ONR and its research community and stakeholders. Accordingly, while this volume is the last you'll see from the Office of Innovation, you will see the Innovation Newsletter transform into a new periodical called Future Force. This new publication will incorporate all elements of research sponsored by ONR that lead to innovations for the future force.

Many of you were with us when Volume I was published in March 2009. While the newsletter could not exist without the fine submissions by the authors of our articles, and while many people contributed to putting the newsletter together (thanks to all!), there are two people most worthy of mention in regards to the communication strategy for the Office of Innovation. Ms. Laura Smith and Ms. Melody Cook both deserve special thanks. Ms. Smith for her strategic vision and ability to make real a dream. And to Ms. Cook who was the true glue that bound these pages together and both tirelessly and very successfully brought each and every single edition from brainstorming to publication and distribution. Lastly to Craig Hughes, who ably served both as Deputy Director and as Acting Director of Innovation – BRAVO ZULU! ■

~ **LARRY SCHUETTE, ONR ACTING DIRECTOR OF RESEARCH**

INNOVATION BEYOND IMAGINATION™

Lawrence C. Schuette, Ph.D.
Acting Director of Research
larry.schuette@navy.mil

Craig A. Hughes
**Acting Deputy Director of
Research**
craig.a.hughes@navy.mil



Office of Naval Research

One Liberty Center
875 North Randolph Street
Suite 1425
Arlington, VA 22203-1995
www.onr.navy.mil/innovate